

Managing Fleet Wide Sensory Data: Lessons Learned in Dealing with Volume, Velocity, Variety, Veracity, Value and Visibility

Preston Johnson¹

¹National Instruments, Austin, Texas, 78759, USA

preston.johnson@ni.com

ABSTRACT

More than ever, asset operators and OEMs are investing in fleetwide monitoring systems. With the roll out of these monitoring systems, huge amounts of sensory data are generated. In a single Gigawatt power plant, asset monitoring systems sort through terabytes of sensory data per week. To contend with the volume and velocity of sensory data, analytics and data management techniques are employed along the life of sensory data from digitization at the asset, to storage in the information technology infrastructure. This paper presents techniques, both promising and fielded, for analytics to manage the volume, velocity, veracity, variety, and value of fleetwide asset monitoring data yielding opportunities for advanced visibility of actionable information.

1. INTRODUCTION

In industrial asset monitoring applications, scientists, engineers, and asset maintainers can collect vast amounts of data every second of every day. Drawing accurate and meaningful conclusions from such a large amount of data is a growing problem, and the term “Big Data” describes this phenomenon. Big Data brings new challenges to prognostics applications in the form of analysis techniques, search and retrieval, data integration or fusion, reporting, and system maintenance (Johnson & Farrell, 2011). All these challenges must be met to keep pace with the experimental growth of asset related data.

Take for example, the Large Hadron Collider at the European Organization for Nuclear Research (CERN), where for every experiment the control and monitoring systems can generate 40 terabytes of data (Badicich & Orci, 2012), (Losito 2011). In Aerospace, for every 30 minutes a jet engine runs, upwards of 10 terabytes of operational data is generated. In a single journey across the

Preston Johnson. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Atlantic Ocean, a four-engine jumbo jet can create 640 terabytes of data. Multiply the single flight by 25,000 flights per day, and we yield an enormous amount of data (Gantz & Reinsel, 2011). This is “Big Data”.

2. HISTORY OF BIG DATA

The technology research firm International Data Corporation (IDC) recently performed a study on digital data, including measurement files (think time waveform recordings), video (think thermal images), music (think ultrasonic), work order reports, and so on. The study estimates that the amount of data available is doubling every two years. In 2011 alone, 1.8 zettabytes (1E21 bytes) of data were created (Hadhazy, 2012), Figure 1. While, our (as in the PHM community) asset monitoring systems may not produce quite this amount of data, just consider the size of the data files we collect from diagnostic visits to our assets. Next consider the impact that low cost automatic data collection systems and sensors can and are having in our ability to continuously monitor and record data from our assets. Even within PHM asset monitoring and prognostics functions, the trends are similar: the amount of data available for predictive analytics is doubling every two years.



Figure 1. Data is collected at a rate that approximately parallels Moore’s law.

The fact that the volume of data is doubling every two years mimics one of the electronics’ most famous laws: Moore’s law. In 1965, Gordon Moore stated that the number of

transistors on an integrated circuit doubled approximately every two years and he expected the trend to continue “for at least 10 years”. Forty-five years later, Moore’s law still influences many aspects of Information Technology (IT) and electronics. Consider that in 1995, 20 petabytes of total hard drive space was manufactured. Today, Google processes more than 24 petabytes of information every single day. Similarly, the cost of storage space for all this data has decreased exponentially from \$228/GB in 1998 to \$0.06/GB in 2010. (Unfortunately, memory sticks at our favorite electronics stores are still a bit more expensive).

Changes, including lower cost of storage and lower cost of data recording devices undoubtedly, fuel the Big Data phenomenon and raise the question, “How do we (the PHM Community) extract meaning from that much information”. Another question might be “What is the value of Big Data”. One intuitive value of more and more data is simply that statistical significance increases. This is certainly the case in data-driven prognostics. Yet, care is required. Consider the gold mine metaphor, where in the mine, only 20 percent of the gold is visible. The remaining 80 percent is in the dirt where it cannot be seen. Mining is required to realize the full value of the contents of the mine. Hence Big Data Analytics and data mining are required to achieve new insights that have never before been seen.

To fully characterize Big Data, consider Figure 2. The challenges of big data are variety, velocity, and volume. These three are often referred to as the three “V”’s of big data. Here we consider three additional V’s, veracity, value, and visibility. Volume is the amount of data as measured in its computer disk or computer memory size. Velocity is the speed at which data is produced, and moved into the computing infrastructure. Veracity is a measure of accuracy or reliability of the data, in other words the validity of data. Variety is both the data structure such as binary files and database tables, and the sources such as vibration, temperature, and maintenance records. Value is the information and business guidance that can be extracted from the data. Last but not least, visibility is the ability to access and view data and its value, regardless of the location of the data within the computing infrastructure.

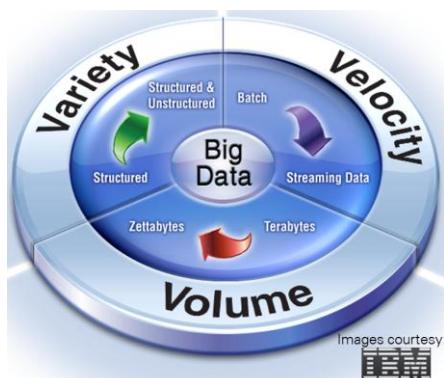


Figure 2. Traditional 3 “V”’s of big data (source: IBM)

3. INDUSTRIAL INSTRUMENTATION, BIG DATA, PROGNOSTICS

The sources of Big Data in the Industrial Asset Monitoring arena are many, Figure 3. The most interesting is data derived, using transducers, from the physical world. In other words, this is analog data captured by instruments and data acquisition systems from a variety of vendors, in a variety of formats. Thus, the PHM community may call it “Big Analog Data” (BAD). BAD is derived from time waveform measurements from vibration, dynamic pressure, thermal images, ultrasonic scans, motor current signatures, and even radio frequency measurements used in the detection of partial discharge or electrical ground faults. Engineers, Scientists, and our plant Maintainers publish this kind of data (BAD) voluminously, in a variety of forms, and many times at high velocities. Along with management and storage of this large amount of data, are the challenges of validation or veracity, deriving value from the data, and giving visibility of data and derived value to the right people at the right time.

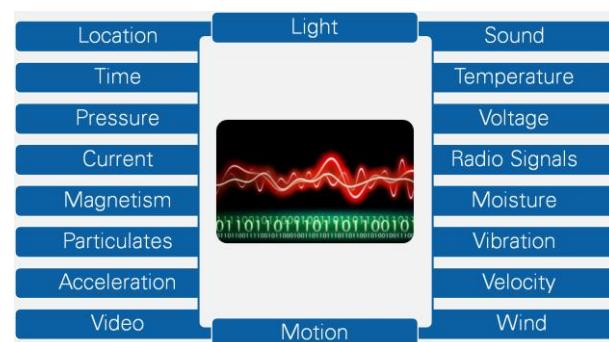


Figure 3. Industrial sources of analog data

As scientists and engineers work to address this “BAD” challenge, an approach is needed that encompasses sensors and actuators, distributed acquisition and analysis nodes (DAANs), and Information Technology (IT) infrastructure for big data analytics, mining and storage. Consider a three-tier solution, Figure 4. Here, it is possible to distribute the work of finding value in big analog data. Figure 4 depicts a three-tier architecture with sensors (and monitored assets) on the left. Measurement hardware or data acquisition systems are in the middle. These devices digitize analog sensory data from a single monitored asset and begin preliminary analysis. The right side of Figure 4 depicts the IT infrastructure employed to store, manage, and analyze sensory data from a fleet of assets.

Two additional terms are introduced here to describe veracity and extraction of value: “In-Motion” and “At-Rest” analytics. With In-Motion analytics, data is analyzed for value in the form of indicative information, in memory, and as close to the source of the data as possible. With At-Rest analytics, data is analyzed in its storage place often

incorporating similarities and differences with collaborative data sources. Both the DAANs and the IT computers perform in-motion analytics, extracting condition indicators. The IT infrastructure, as it assembles sensory and other data from multiple sources, also performs at-rest analytics utilizing data-driven prognostic algorithms to identify patterns and fault signatures.

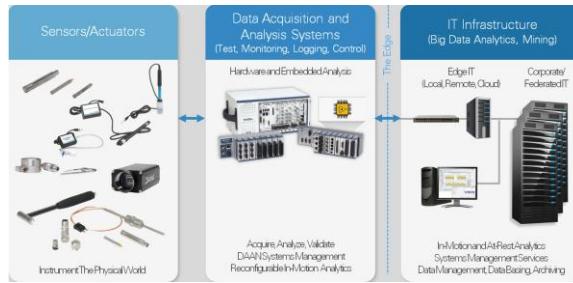


Figure 4. A three-tier solution to the “Big Analog Data” challenge.

Let's look closer at in-motion analytics close to the sensor. For example, adding a smart chip such as a Field Programmable Gate Array (FPGA) or a processor to an analog sensor allows the sensor to reduce the raw analog data to condition indicating features of the time waveform. However, it is also possible to add “smart” data recorders to the traditional analog sensors installed today. Both the smart sensor and the smart recorder are able to implement a decision based data recording technique, Figure 5. Here, analog sensory time waveform data is continuously analyzed for changes. Only when an indication of change within the asset is present in the sensory data (or on a time basis for periodicity) is the data recorded and forwarded upstream in the three-tier architecture. Further, the sensory data might be reduced using in-motion analytics to a set of condition indicators or features, leaving the raw time waveform stored locally or discarded. The filtering process of looking for changes and reducing data to condition indicators plays a big role in managing volume, velocity, veracity, and value.

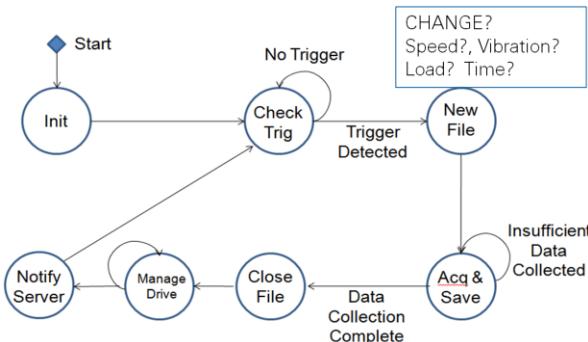


Figure 5. Decision based data recording state diagram

Whether, we have the ability to perform analysis in-motion at the sensor, at the DAAN or at-rest in the IT Infrastructure,

we are fortunate to have a number of analytical tools at our disposal for finding value in the data. The scientific fields of condition monitoring and prognostics offer a number of analytical tools for reducing data to condition indicators and for finding trends in the analytical results, Table 1, Figure 6. Condition indicating analytics range from vibration level measurements, temperature trends, to envelope spectrum for roller bearing degradation and so on. With condition indicating analytics, we can discover increased impacting in roller element bearings, teeth cracking in gearboxes, rotor bar degradation in induction motors and generators, and so on. Condition indicators, coupled with trending and alarming, give the asset owner / operator a first alert that degradation is occurring within the asset.

Table 1. Condition indicating analytics

Graphic	Signal Characteristic	Analysis Methods	Machine Example
	Narrow frequency band lasting for a long time	Frequency Analysis Fourier Transform Power Spectrum	Unbalance in a single speed machine
	Narrow frequency band with harmonics lasting for a long time	Quefrency Cepstrum	Damaged bearing in a machine with roller element bearings
	Time varying frequency band	Time-frequency analysis Order analysis	Unbalance in a variable speed pump
	Wide frequency band signal lasting for a short time	Wavelet analysis AR Modeling	Low speed machine with compressor valve impacts
	Narrow frequency band signal lasting for a short time	Wavelet Analysis	Electrical motor driven machine with rub and knock noise.

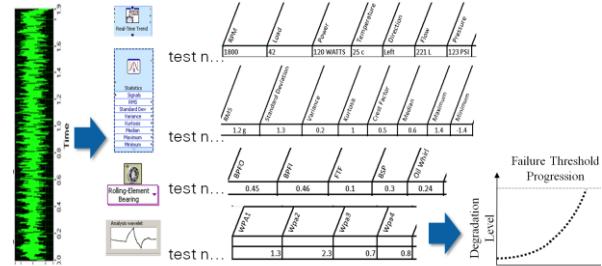


Figure 6. Reducing sensory data to condition indicators

Within the PHM community, the use of multiple condition indicators in concert, and an extensive history of actual condition indicators, data driven prognostics is made possible. Prognostic analytics include clustering, statistical pattern recognition, logistic regression, support vector machine, neural networks and so on. These are similar mathematics used in big data sciences, a growing profession and industry sector. Together, these two classes of analytics (condition indicators and prognostics) provide the foundation for finding value in big analog data. Long term, these tools are building the foundation for automating diagnostics, and prognostics. With the automation of diagnostics and prognostics, business decisions can be

enhanced with automatically generated advisories for maintenance, operations, and finance.

The condition indicators themselves do not necessarily yield a root cause for the degradation, nor does the condition indicator tell us when we can expect the asset to fail to perform its function. Prognostic analytics are employed to help deduce the why and when of asset degradation and failure, Figure 7.

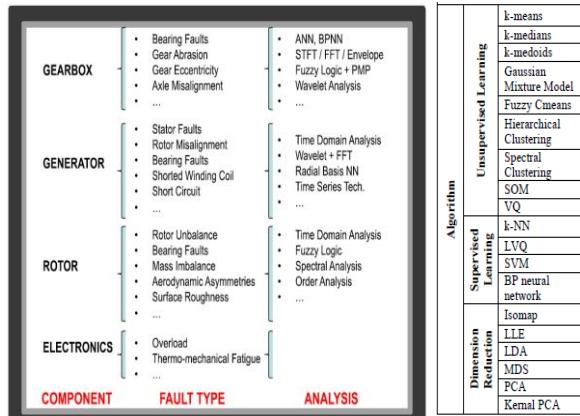


Figure 7. Prognostic analytics for finding patterns

Prognostic algorithms allow for the combination and collaboration of condition indicators within an asset (bearing, gear, shaft, oil particle, temperature, load, speed) as well as across similar assets. This combination of condition indicators forms a pattern of healthy asset operation, or a specific degradation pattern. In practice, a baseline of healthy condition indicators is obtained during commissioning of an asset, or after repair and maintenance of an asset. With an available healthy or normal operation pattern, analytical tools including statistical pattern recognition can be used to determine electrical, mechanical, or structural degradation levels of an asset, Figure 8. These tools compare real-time sensory data in-motion to patterns looking for deviations or anomalies.

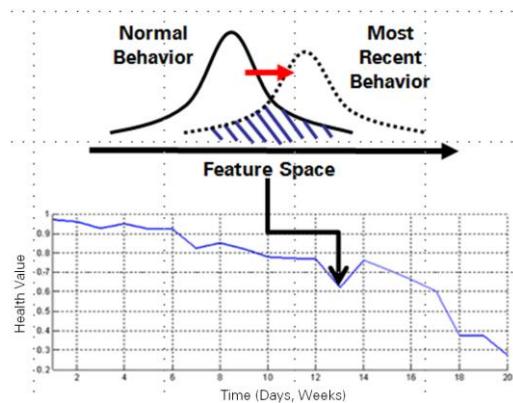


Figure 8. Asset degradation using statistical pattern analysis

The normal and fault patterns are further extended, by further segregating these patterns into operating conditions when speeds, loads, and environment are included. The combination of patterns at a plant or enterprise level, is made possible when similar assets are viewed together, enhancing the pattern formation. For example, machine learning algorithms are able to cluster combinations of condition indicators from similar assets, thereby creating patterns of normal or fault asset behavior. Prognostic algorithms then use these patterns, or fault signatures, to match current asset condition indicators to a specific fault signature (with in-motion analytics).

On another note, as condition indicators are narrowed in number to the best indicators of specific failure modes, a smaller set of sensors and analytics may be used to detect and predict specific failure modes. These reduced sensory measurements and analytics can then be performed on sensory data in-motion on the (embedded) DAAN, comparing a single vector of condition indicators to specific fault patterns.

As the normal operational pattern “drifts” towards a specific fault signature pattern, the rate of “drift” combined with human expert knowledge to form a basis for automatic advisory generation and prediction of the point in time when the asset fails to perform its function. This is particularly true at the information technology (IT) level, when future operating conditions are known based on planned equipment operations. Knowledge of a future operating condition allows focus on data-driven patterns from historical and specific expected operating conditions. Trends derived from historical specific operating conditions, improve confidence in the expected performance and health of specific equipment in planned operating conditions. At the plant or even enterprise level, the fusion of operational and equipment data builds a foundation for and confidence in the data-driven predictions.

To summarize, there are many physical phenomenon to measure within a fleet of assets. This creates the big data problem of the analog kind. By using in-motion and at-rest analytics, the six V’s of big analog data are addressed. Analytics that calculate condition indicators, derive patterns of condition indicators, and compare real-time condition indicators to normal and faulty patterns are core to addressing the challenge of big analog data. This challenge of big analog data is deriving value and visibility while managing volume, velocity, veracity, and variety.

4. INFORMATION TECHNOLOGIES

In addition to sensory data, condition indicators, and asset operational patterns, we (the PHM community) often add other data which may be unstructured in nature. Work order reports, typed textual descriptions, and diagnostic technical

exams add to our big analog data, extending our view of the health of assets. To support big analog data storage and analytics as well as varied documentation, consideration and collaboration with our colleagues in Information Technology (IT) is a must.

Part of our challenge with big analog data and the varied documentation formats, is the data does not fit easily into standard relational databases. As a comparison, neither does the vast information available on the world-wide web. Out of Google's work to "index" the web, came an underlying file system, Apache Hadoop, which supports unstructured data or data that is stored in files rather than a relational database, Figure 9. These files can include binary and ASCII formats of condition indicators and time waveforms. Our unstructured data files also include asset technical exam documentation. There are many common formats used for big analog data including UFF58, Comtrade, and .mat. In the case study presented later, the file structure named Technical Data Management Streaming (TDMS) is used for storing time waveforms and condition indicators. The Apache Hadoop File System (HDFS) helps to manage these non relational database items. The HDFS is a massively scalable storage and batch data processing system. It provides an integrated storage and processing fabric that scales horizontally with commodity hardware and provides fault tolerance through software. Hadoop also includes concepts for distributing analytics to the data, to avoid bandwidth issues of moving the at-rest data (Bisciglia, 2009).

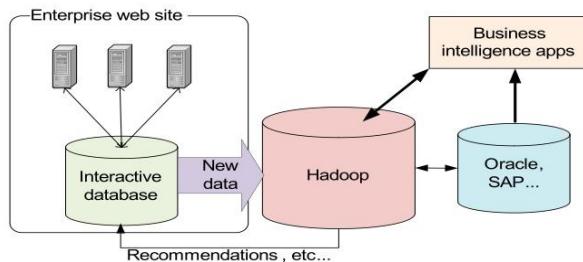


Figure 9. High level overview of Hadoop file system within IT architecture (source: Cloudera)

Several information technologies suppliers take the concept further by industrializing HDFS and improving the programming tools used to mine and analyze the data in a combination of Hadoop and relational stores. International Business Machines (IBM) for example, not only hardens the IT infrastructure with their "PureFlex" enterprise computing systems, IBM also adds InfoSphere Streams for in-motion analytics and InfoSphere BigInsights for at-rest analytics, Figure 10. These architectures and analytic tools promise an ability to quickly garner value of our variety, velocity and volume of Big Analog Data and unstructured documentation (Franklin, 2012).

The convergence of pervasive sensory data sources, new information technologies, growing information stores and a reduction in the overall cost and time needed for analysis has helped big data and specifically our industrial big analog data cross the chasm from innovation to early adoption. Big data is still an early-stage technology, but expect that over the next 18 months it will break double digits on project adoption basis. (Rogers, 2011).

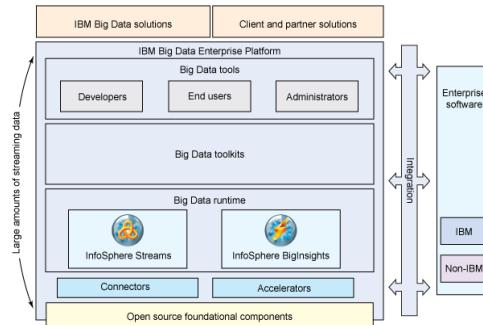


Figure 10. IBM's platform and vision for big data (source IBM DeveloperWorks)

So, if we can combine big analog data, in-motion and at-rest analytics of the condition indicating and prognostics kind, with expanded information technologies; perhaps it becomes possible to create smart monitoring and diagnostics, or even cloud based prognostics. The Center for Intelligent Maintenance systems projects a future where multiple end users will submit their asset data and condition indicators to a cloud resource (IMS, 2012) Here, analytical collaboration occurs to build and leverage fault signatures, degradation patterns, along with prognostic analytics to advise us on the current and future health of our assets, Figure 11.

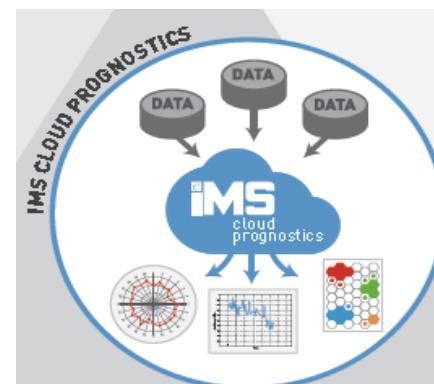


Figure 11. Center for Intelligent Maintenance Systems Cloud Prognostics Vision (source: IMS Center)

Given that Moore's law of big data is a true observation, then the doubling of data every two years demands that these information technologies will mature and become more pervasive. The field of prognostics will benefit from

the collaboration that comes with a wide net of assets, sensory data, and condition indicators derived from the sensory data. The combination of prognostics and data science technologies with information systems technologies is already yielding solutions for the volume, velocity, veracity, variety, value, and visibility of the fleetwide monitoring big analog data challenge.

5. CASE STUDY

In power generation, the above mentioned technologies are coming together to solve fleetwide asset monitoring data and information challenges. The Electrical Power Research Institute (EPRI) continues to sponsor a fleet wide asset monitoring project within a special working group, the Fleetwide Monitoring Interest Group (FWMIG) (Hollingshaus, 2011). This program aims to articulate a condition based maintenance and prognostics solution for its power generation members. The applications framework leverages data available within power generation plants, a fault signature database, and traditional monitoring and analysis techniques for rotating machinery.

Duke Energy, an EPRI member, is already deploying hundreds of new low cost “smart” data acquisition and analysis nodes (DAAN) within several power generation plants (Cook, 2013). These DAANs use traditional piezoelectric dual mode accelerometers with temperature sensing elements to monitor for changes in balance of plant equipment that supports turbine generators, Figure 12.

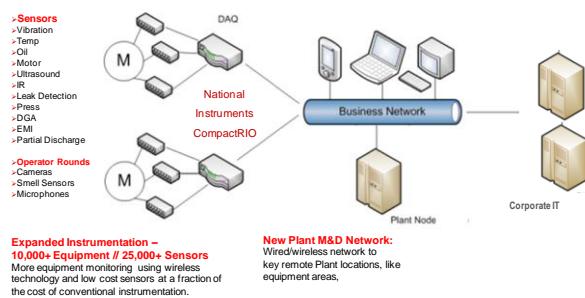


Figure 12. Duke Energy architecture for data acquisition and analysis nodes.

In the late 1990s, Duke Energy began its fleetwide monitoring program using commercial handheld instruments for vibration, thermography, ultrasonic, motor current, and oil analysis. Today, Duke Energy machinery health subject matter experts spend 80 percent of their time with these hand held instruments simply collecting sensory data.

Beginning in 2012, Duke Energy began to automate data collection with flexible DAANs, thereby reducing the labor costs and sparse periodicity associated with manual analog data collections. With the new DAANs in place, these same

subject matter experts will be able to spend 80 percent of their time analyzing sensory data and planning maintenance actions. While the core initial motivation and return on investment at Duke Energy is employee utilization, the opportunity for prognostics, especially data driven, is tremendous as vibration, temperature, and oil analysis analog data now stream at regular intervals into the Duke Energy IT infrastructure, Figure 13.

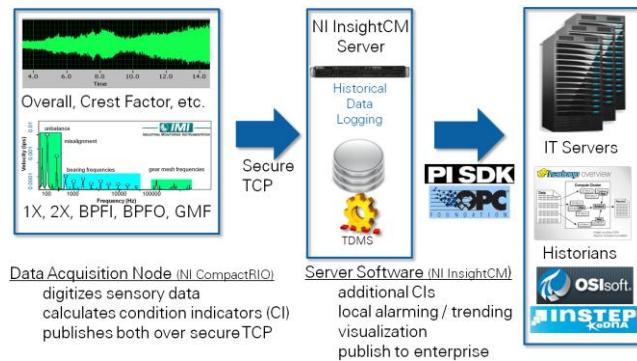


Figure 13. Big analog data sensory data flow

To accomplish the high level architectures, Duke Energy is working with EPRI and condition monitoring vendors to develop and implement a big analog data system for fleetwide asset monitoring that manages the six “V” challenges of big data. As shown earlier in Figure 5, and in Figure 13, the DAAN works to address volume, velocity, veracity, variety, and value. Using an event base local recording structure, Figure 5, sensory data is filtered to just data that is periodic or has a change. This filtering helps address volume. Using a store and forward communications scheme, data is transferred at the bandwidth allowed on the network. By storing and forwarding, the velocity of data is controlled by network administration tools. The DAAN also checks sensor value validity by using range checking and open/short cabling issues. This sensor value check helps address veracity. Lastly, the DAAN labels all data with sensory data type, measurement characteristics, and equipment hierarchy down to the component where the sensor is attached. The labeling tasks helps address the variety of the various analog measurements made by the DAAN.

To support the new volume, velocity, and variety of data coming from the newly deployed DAANs, Duke Energy has formed an IT task force to develop a big analog data strategy. The goal of the task force is to maximize value and visibility in particular with respect to equipment maintenance, availability and reliability. The current organization of data analytics orchestrated by Duke Energy IT, EPRI, and vendors is show in Figure 14. Value and Visibility at Duke Energy are determined at the monitoring and diagnostics center in Charlotte, NC. Here all condition indicators and operational process parameters are recorded in OSIsoft PI™'s historian for advanced pattern recognition

and anomaly detection by Instep Software's PRiSM™ predictive analytics tools.

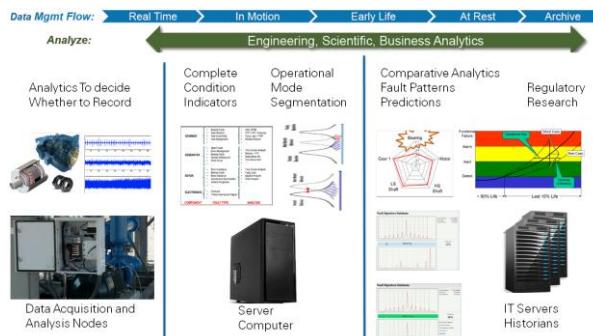


Figure 14. Analytics flow in big analog data applications

While the condition indicators are published to enterprise historians, the technical exam data including vibration time waveforms, stored in TDMS format, remains at the plant server level. This allows subject matter experts to access and analyze the analog sensory data using common graphics and analysis techniques associated with the particular technology. For example, vibration time waveforms are analyzed with frequency spectra, in the order domain, using harmonic, sideband cursors, and waterfall displays. The vibration analytical tools also provide trends and alarms at the local plant level for harmonics of rotational speed or order analysis, as well as trending of all condition indicators calculated at the DAAN or the plant server computer level.

However, time waveform data is big data, and the volume needs management at the plant level. Once condition indicators are extracted and published to the OSIsoft PI™ historian, some of the time waveform data can be discarded. An aging strategy is implemented that removes all time waveform data, after five days with the exception of those time waveforms most close to peak power demand times of day, 8:00 AM, Noon, and 4:00 PM. In addition, any time waveform that was recorded due to a measurement value alarm is preserved. Subject matter experts can also mark specific data files for preservation as the need arises.

As condition indicators are analyzed in the historian, user notes regarding equipment, maintenance records, best practices, and recommended actions are also assembled from various data sources and locations within the Duke Energy information technology infrastructure (Hesler, 2010). The challenge lies in assembling, storing, and retrieving information both from fleetwide asset monitoring and also operating parameters, maintenance activities, and equipment component health. To address the challenge, Duke Energy has deployed EPRI's PlantView® software platform for managing power plant assets and developing condition status reports on plant equipment, Figure 15.



Figure 15. PlantView® health report matrix, image courtesy of Power Vision, Inc.

The PlantView software provides applications for entering, storing, and viewing information about plant operating parameters, maintenance activities, and equipment health. The status of equipment is kept in an integrated database. Visibility is provided thru a series of web services applications allowing users to access information from user customizable web portals. Duke Energy now has over 10,000 internal users benefiting from the PlantView web portals.

At Duke Energy, this is an obvious case where the opportunity for prognostics and IT come together to mine big analog data for the benefit of asset owners, asset operators, and the evolution of prognostics. Beginning with the DAAN, condition indicators extracted from monitored equipment, are supplemented with additional condition indicators at the plant server computer. This is the same computer that manages the DAANs. Subsequent to publishing the condition indicators to the enterprise historian, the advanced pattern recognition software begins comparison of current condition indicators to baselines for the specific operating condition. A web interface is provided for systems users and business owners to see both power output from generating units, as well as any equipment or process problems that may need addressing. The web interface, PlantView, brings the value and visibility of operations data to those responsible for making business decisions.

6. CONCLUSION

Big data, especially of the analog kind, can and does present challenges. Fortunately, information technology is evolving as quickly as the volume of data grows. Both in-motion and at-rest analytics are working to make sense of big analog data. The growing deployment of a wide range of sensors across a wide net of assets promises to accelerate the success and science of prognostic applications for monitoring fleets of assets.

REFERENCES

Bisciglia, C. (2009). 5 Common Questions About Apache Hadoop. *Cloudera Blog*, 14 May 2009

<http://blog.cloudera.com/blog/2009/05/5-common-questions-about-hadoop/>

Bradicich, T. & Orci, S. (2012). Moore's Law of Big Data *National Instruments Instrumentation News*. December 2012. Web. <http://zone.ni.com/devzone/cda/pub/p/id/1649>

Center for Intelligent Maintenance Systems (IMS), (2012). IMS Center brochure. *IMS Center website*. 5 Aug 2012. http://www.imscenter.net/Resources/brochure_2012_re_d_final.pdf

Cook, B. (2013). Deploying smart maintenance and diagnostics for electrical power generation. *NIWeek 2013*. 7 Aug 2012. <https://decibel.ni.com/content/docs/DOC-30892>

Franklin, C. (2012). Big Data as part of an enterprise data strategy. *Tamgroup Blog*. 19 March, 2012 <http://www.tamgroup.com/blog/bid/118927/Big-Data-as-part-of-an-enterprise-data-strategy>

Gantz, J., & Reinsel, D. (2011). Extracting value from chaos. *EMC Corporation website*. June 2011. Web.

Hadhazy, A. (2012). Zettabytes now needed to describe global data overhead. *Live Science*. 4 May 2010 Web.

Hessler, S. & Noce, G. "New web applications in EPRI's PLantView software offer Progress Energy enhanced capabilities for using plant data", publication 1021286, Electrical Power Research Institute, Palo Alto, California, USA

Hollingshaus, B. (2011). Program 69: Maintenance management and technology. *Electrical Power Research Institute Descriptions of Past Research*. Catalog number 1022681, May. 2011 www.epri.com

Johnson, P. & Douglas F. (2011). The impact of rapidly changing computing technologies on prognostic asset management applications. *MFPT Newsletter*. November 2011. Web. <http://www.mfpt.org/Newsletters/1111/Johnson.htm>

Losito, R. (2011). World's largest particle accelerator" *National Instruments Case Study*. 2011. Web. <http://sine.ni.com/cs/app/doc/p/id/cs-10795>

Rogers, S. (2011). Big data is scaling BI and analytics. *Information Management*. 1 Sep 2011.

BIOGRAPHIES

Preston Johnson is the Principal Sales Engineer for Condition Monitoring Systems at National Instruments (NI) in Austin, Texas. He has worked for National Instruments for over 27 years in roles of Field Sales, Sales Management, Automation Business Development, Sound and Vibration Segment Manager, Platform Manager for Condition Monitoring Systems and Global Program Manager for Asset Monitoring Systems. In his current role as Principal Sales Engineer, Preston works with NI OEM and End User customers to deploy fleetwide asset monitoring systems that lower operation costs, improve machinery reliability, and ultimately increase revenue. His interests lie in embedded signal processing, data acquisition systems and architectures, and prognostics. He earned his BSEE in Electrical Engineering and Computer Science from Vanderbilt University in 1985 and his MBA in Information Systems from the University of Texas in 1987. Preston is experienced in project management and holds a Category III vibration analyst certificate.